



AFRL-RI-RS-TR-2017-043

## ASSISTED PERCEPTION, PLANNING AND CONTROL FOR REMOTE MOBILITY AND DEXTEROUS MANIPULATION

---

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

APRIL 2017

FINAL TECHNICAL REPORT

**APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED**

STINFO COPY

AIR FORCE RESEARCH LABORATORY  
INFORMATION DIRECTORATE

## **NOTICE AND SIGNATURE PAGE**

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 88<sup>th</sup> ABW, Wright-Patterson AFB Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2017-043 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /  
MICHAEL COREY  
Work Unit Manager

/ S /  
JULIE BRICHACEK  
Chief, Information Systems Division  
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

# REPORT DOCUMENTATION PAGE

**Form Approved  
OMB No. 0704-0188**

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

**PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

<b>1. REPORT DATE (DD-MM-YYYY)</b> <b>APRIL 2017</b>		<b>2. REPORT TYPE</b> <b>FINAL TECHNICAL REPORT</b>		<b>3. DATES COVERED (From - To)</b> <b>SEP 2012 – AUG 2016</b>	
<b>4. TITLE AND SUBTITLE</b>  <b>ASSISTED PERCEPTION, PLANNING AND CONTROL FOR REMOTE MOBILITY AND DEXTEROUS MANIPULATION</b>		<b>5a. CONTRACT NUMBER</b>			
		<b>5b. GRANT NUMBER</b> <b>FA8750-12-1-0321</b>			
		<b>5c. PROGRAM ELEMENT NUMBER</b> <b>62702E</b>			
<b>6. AUTHOR(S)</b>  Russell Tedrake, Michael Corey		<b>5d. PROJECT NUMBER</b> <b>ROBO</b>			
		<b>5e. TASK NUMBER</b> <b>PR</b>			
		<b>5f. WORK UNIT NUMBER</b> <b>04</b>			
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> Massachusetts Institute of Technology 77 Massachusetts Avenue Cambridge, MA 02139				<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>	
<b>9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> Air Force Research Laboratory/RISC 525 Brooks Road Rome NY 13441-4505				<b>10. SPONSOR/MONITOR'S ACRONYM(S)</b> <b>AFRL/RI</b>	
				<b>11. SPONSOR/MONITOR'S REPORT NUMBER</b> <b>AFRL-RI-RS-TR-2017-043</b>	
<b>12. DISTRIBUTION AVAILABILITY STATEMENT</b>  Approved for Public Release; Distribution Unlimited. PA# 88ABW-2017-1248 Date Cleared: 22 Mar 2017					
<b>13. SUPPLEMENTARY NOTES</b>					
<b>14. ABSTRACT</b> <p>This report describes the methods, procedures and results of developing software infrastructure for allowing remote control of a bipedal human-scale Atlas robot. This work focuses on allowing remote operations in severely communications restricted environments for search and rescue. The research relied heavily on abstraction methods for issuing commands, compression and processing of sensor data, and ensuring the safety of the robot during execution. The DARPA Robotics Challenge provided a live demonstration to showcase the robots performance in a simulated search and rescue environment. The abstraction techniques developed from both a control and terrain estimation viewpoint, were successfully able to allow autonomous operation. These techniques substantially reduced the amount of communication necessary between the human operator and the robot allowing them to operate with substantial delays in communication as well as complete losses of communication. The software developed was successfully able to control the robot to complete a significant majority of tasks to drive a vehicle, dismount the vehicle, open a door, manipulate various objects and ultimately walk through uneven surfaces and climb stairs. MIT placed sixth in the finals which was an impressive finish given that their best run required them to operate without the use of one arm which was damaged in a fall.</p>					
<b>15. SUBJECT TERMS</b> Robotic Control, Robots, Degraded Coms					
<b>16. SECURITY CLASSIFICATION OF:</b>		<b>17. LIMITATION OF ABSTRACT</b>  UU	<b>18. NUMBER OF PAGES</b>  20	<b>19a. NAME OF RESPONSIBLE PERSON</b> <b>MICHAEL COREY</b>	
a. REPORT U	b. ABSTRACT U			c. THIS PAGE U	<b>19b. TELEPHONE NUMBER (Include area code)</b> <b>N/A</b>

## Table of Contents

List of Figures.....	ii
1.0 SUMMARY.....	1
2.0 INTRODUCTION.....	1
3.0 METHODS, ASSUMPTIONS, AND PROCEDURES .....	2
3.1 TERRAIN PERCEPTION .....	2
3.2 WHOLE-BODY STATE ESTIMATION .....	3
3.3 FOOTSTEP PLANNING .....	4
3.4 WHOLE-BODY DYNAMIC MOTION PLANNING.....	5
3.5 WHOLE-BODY FEEDBACK CONTROL .....	6
3.6 MANIPULATIONS.....	6
4.0 RESULTS AND DISCUSSION .....	7
4.1 DRC FINALS COMPETITION .....	7
5.0 CONCLUSIONS.....	12
6.0 REFERENCES.....	13
LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS.....	15

## List of Figures

Figure 1 - An overview of the transformation from images to point clouds and surfaces.....	2
Figure 2 - Demonstration of the robot's dynamic balancing .....	4
Figure 3 - Bipedal multi-dimensional motion estimation planning .....	5
Figure 4 - Balance feedback demonstration of the Atlas robot.....	6
Figure 5 - User view of Atlas robots procedural manipulation planning .....	7
Figure 6 - MIT's Atlas Robot performing driving task at DRC Finals .....	8
Figure 7 - MIT's Atlas robot failing to successfully dismount the Polaris Ranger.....	9
Figure 8 - Atlas Robot successfully opening a door and crossing the threshold. ....	9
Figure 9 - A remote view of the valve task as seen by the human operator at base station .....	10
Figure 10 - Traversal of the uneven terrain at the DRC finals.....	11
Figure 11 - The robot successfully traversing the stairs obstacle .....	12

## **1.0 SUMMARY**

This report describes the methods, procedures and results of developing software infrastructure for allowing remote control of a bipedal human-scale Atlas robot. This work focuses on allowing remote operations in severely communications restricted environments for search and rescue. The research relied heavily on abstraction methods for issuing commands, compression and processing of sensor data, and ensuring the safety of the robot during execution. The DARPA Robotics Challenge provided a live demonstration to showcase the robots performance in a simulated search and rescue environment.

The abstraction techniques developed from both a control and terrain estimation viewpoint, were successfully able to allow autonomous operation. These techniques substantially reduced the amount of communication necessary between the human operator and the robot allowing them to operate with substantial delays in communication as well as complete losses of communication.

The software developed was successfully able to control the robot to complete a significant majority of tasks to drive a vehicle, dismount the vehicle, open a door, manipulate various objects and ultimately walk through uneven surfaces and climb stairs. MIT placed sixth in the finals which was an impressive finish given that their best run required them to operate without the use of one arm which was damaged in a fall.

## **2.0 INTRODUCTION**

Fueled by a sense of urgency and purpose provided by the DARPA Robotics Challenge (DRC), over the course of this program MIT has made dramatic progress in remote-assisted dexterous manipulation (and legged locomotion). Their work earned them a fourth place finish in the 2013 DRC Trials, and a sixth place finish in the 2015 DRC Finals. Those numbers only tell a part of the story. Beyond providing robust implementations of known best algorithms, the team consistently pushed the frontiers of research. Our system was widely considered one of the most autonomous of the capabilities, and the research papers published about the technology have won Best Paper awards (or Best Paper Finalist) at the annual Humanoid Robotics Conference every year since they began publishing in the domain.

The research developed novel algorithms, and also provided a detailed post-mortem analysis of the performance in the competitions. Moreover, throughout the duration of the program, the performers have continued to push nearly all of the developed code to open-source on [github.com](https://github.com), have used these in a robotics MOOC (massive open online course), and continue to support the growing community of users. At the DRC Finals, DARPA recognized the team with the Seth Teller Memorial Award for Contributions to Open-Source Software.

The following pages provide a high-level summary of the advances, with pointers to research papers and reports with the specific algorithmic details.

## 3.0 METHODS, ASSUMPTIONS, AND PROCEDURES

### 3.1 TERRAIN PERCEPTION

In the DRC finals, they fielded a terrain perception system that made primary use of LIDAR. They produced a general-purpose height map estimation functionality which allows the robot to walk on arbitrary terrain. Additionally, the research introduced specialized perception algorithms for fitting terrain primitives — especially cinder blocks and stairs. In addition to single primitives, composite groups of geometry are also supported. These groupings can be used both to support the stairs use case as well as various configurations of stacked and tilted cinder blocks as seen in the DRC.




Figure 1 - An overview of the transformation from images to point clouds and surfaces

Since the geometric primitives are aligned using noisy sensor data, a method to continually adjust and reevaluate their position and orientation was implemented. This ensures that as progress is made across the terrain, new information is incorporated to maintain the accuracy of the perceived map.

When combined with robust online replanning of footsteps, this enabled the performer to achieve a very high success rate in our attempts to traverse uneven terrain environments with the Atlas bipedal robot.

The LIDAR provided clean and high-resolution data even outdoors, but at a very low rate (this was the primary reason why our robots moved slowly during the competition). Our algorithms would wait for an entire revolution of the LIDAR before attempting to walk, and they opted to spin the laser relatively slowly in order to obtain a dense point cloud. In the lab, the team demonstrated a much more

aggressive approach based solely on stereo vision [12], which enabled continuous walking over the same terrain.

MIT additionally developed an approach for humanoid robots to perform mobility reliably and efficiently over rugged irregular terrain. They researched a method for a robot to utilize only passive stereo imagery to plan footsteps to continuously walk over challenging terrain. The stereoscopic cameras were capable of building a three dimensional model of the terrain. This three dimensional model was met with a quadratic optimization plan which developed an optimal plan over step positions allowing continuous motion while minimizing latency.

The stereoscopic imagery was not used in the Finals due to the inability to test the algorithm under sufficiently varied (outdoor) lighting conditions.

### **3.2 WHOLE-BODY STATE ESTIMATION**

The whole body state estimator Pronto [7] was developed by building upon an existing state estimator developed by MIT that has been proven on unmanned aerial vehicles (UAVs). The underlying algorithm is based on an Extended Kalman Filter (EKF) that simultaneously estimates robot state and sensor biases.

The filter developed provided a probabilistic fusion of sensor data from many modalities to produce a single consistent position estimation for a walking humanoid. Given a prior map using a Gaussian particle filter, the LIDAR based system is able to provide a drift-free alignment resulting in reliable localization of the robot.


The module can also be bootstrapped to providing an on-the-fly map of the system and is capable of performing very robustly in challenging field situations. Additionally, their estimation hierarchy utilizes a two-tier infrastructure allowing both low level control of the robot while preserving registration of the robot and objects in the robots vicinity.

In addition to the information provided by the Inertial Measurement Unit (IMU), Pronto is able to fuse data from a variety of sensors that are commonly found on bipedal robots. This includes laser range finders, stereo cameras, joint encoders, and force-torque contact sensors. Moreover, the state estimator has been extended to leverage insights into the nature of contact to provide estimates that surpass the accuracy of uninformed methods. An example of this is knowledge of when the robot's foot is in contact with the ground and exploiting that information by constraining the kinematic pose being estimated.

The research has continued to develop the state estimator using a simple dynamic model and a complete kinematic model (including the effects of torque-deflection, etc) following the approach in the paper on "Drift-free state estimation"[7]. The algorithm has been optimized to run at rates that exceed 333 Hz on commodity hardware to ensure that reliable estimates are available at rates that enable fast and dynamic responses.

### 3.3 FOOTSTEP PLANNING

Research on explicitly addressing the combinatorial aspects of footstep planning by formulating mixed-integer convex optimization has been one of the most-cited results from this effort [8]. The team has continued to refine and optimize the speed of this algorithm, getting it now to the point where it is fast enough to support dynamic step recovery. If unexpected disturbances or conditions would cause the robot to fall, a reactive step recovery algorithm has been implemented based on the concept of the Instantaneous Capture Point (ICP). The ICP is continually evaluated so that if a fall is about to occur, a stabilizing foothold is available at all times. This allows the robot to switch to a recovery plan that quickly places a foot at the ICP to regain stability from a situation that would otherwise result in a loss of balance. The image below shows the robot reacting to being pushed over by placing its right foot behind itself to catch its fall.



*Figure 2 - Demonstration of the robot's dynamic balancing*

Although such dynamic footstep plans were demonstrated during the DRC finals where the robot was able to save itself from a fall during the uneven terrain task, this is a relatively new advancement and there is still much work to be done in this area. Some of the possibilities include adding additional constraints to the stabilizing footholds to account for known structure in the terrain and using other parts of the robot such as hands and elbows to aid in recovery.

The team continues pushing in research in this direction. In addition, they worked through versions of this planner that support quadruped models and aerial phases where a primary challenge is finding a

suitable convex relaxation for angular momentum. Significant progress in this direction was reported in [21].

This research primarily focused on dynamic potion planning for legged robots specifically in terms of joint movement. This research focused on planning which limb should touch a surface, which order this should happen in and how much force should be applied when impacting the surface. The research leveraged centroidal dynamics to determine the joint angle trajectories in terms of placement.

Optimization was performed through a mixed integer quadratic program. This program took into account center of mass as well as angular momentum trajectories. The research resulted in novel approaches for bipedal gait selection over irregular terrain, trajectory optimization formulation for floating-base systems, and a planning methodology integrating extremely dynamic motions in a diverse live environment.

### 3.4 WHOLE-BODY DYNAMIC MOTION PLANNING

Where previously the connection between the footstep planner and whole-body motion planner was relatively simple (involving only footstep regions), the footstep planner extensions described above mean that it is solving more and more of (a convex approximation of) the full body motion planning problem. Considerable effort was devoted towards exploring this boundary and highly dynamic whole-body plans were found more efficiently and more reliably given a two-step approach [20,22].

A recent discovery made by members of the research team revealed that there is a closed form solution for the optimal controller for stabilizing Zero-Moment Point (ZMP) trajectories. A highly optimized version of this algorithm has been implemented in Drake. Using this technique, solving for optimal ZMP stabilizing controllers can be performed in sub-millisecond computation time [11].

Since this can now be performed in real-time, continuous re-evaluation of the walking controller allows deviations in the planned footsteps to be corrected for by always using the most recently computed gait. This not only alleviates errors that would otherwise accumulate as the plan is executed, but also increases stability and robustness to the point where the robot is able to successfully complete walking tasks where footprint tracking is intentionally degraded by up to a factor of 10. Prior to the implementation of this algorithm, such a tracking error would cause the robot to immediately fall.




Figure 3 - Bipedal multi-dimensional motion estimation planning

The image above shows footstep tracking performance during real-time recomputation of the walking plan. Although the tracking performance was intentionally degraded by an order of magnitude, the final position is still reached by continuously recomputing the controller online.

### 3.5 WHOLE-BODY FEEDBACK CONTROL

One of the early results of this program was the development of an optimization-based feedback control approach that continued to perform well for complex control tasks on a humanoid throughout the competition [2]. By the end of the program, code was significantly optimized. This resulted in a control rate up to nearly 1 kHz by separating the “plan evaluation” into a separate process from the main quadratic-program-based controller. This gave a significant performance increase in terms of robustness, and dramatically reduced (essentially eliminated) the number of times that our “Fast-QP” algorithm had to fall back on the commercial QP solver.

The execution speed of the controller is critical in enabling dynamic and reactive behavior. Moreover, since the controller is now able to run at rates that far exceed the natural timescale of the underlying mechanical system, it can make small adjustments that allow for delicate balancing tasks. The image on the right shows the robot demonstrating this in a task that involves balancing on one foot. A slower controller would have a difficult time producing the fine-grained corrections quickly enough to maintain balance in this situation likely resulting in a fall or significant loss of balance.

The current formulation is carefully integrated with “ZMP-style” walking algorithms for gait planning on flat terrain with a nearly constant center of mass height. Current experiments focus now on stabilizing the more diverse plans that come out of the whole-body planning code, and testing them on the real robot.

Recently, the QP controller developed for the Atlas robot has been applied to the control of the Boston Dynamics Little Dog quadruped robot. This work has shown that the control concepts and optimizations developed for Atlas can easily be extended to the quadruped domain.

### 3.6 MANIPULATIONS

Although it was anticipated that dexterous manipulation would be a major focus of research over the duration of the program, a number of reasons changed this focus. First and foremost, the original robot hands that were capable of dexterous manipulation turned out to be too fragile to be used reliably in competition. Instead, the team and most of other teams elected to use the three-fingered gripped from Robotiq – effectively a large but robust robot claw. As a consequence, DARPA reduced the scope and complexity of the manipulation tasks.

One area that did require considerable research attention was perception for manipulation, and the low-latency remote operation of the manipulation system. They developed a graphical user interface, Director, which not only served as the tool for piloting our robot, but also as a tool for rapidly



Figure 4 - Balance feedback demonstration of the  
Atlas robot

prototyping perception algorithms [17]. Written in python and VTK, it provided hugely valuable libraries for writing algorithms for and visualizing point cloud data.




Figure 5 - User view of Atlas robots procedural manipulation planning

By the end, the robot was mostly autonomous, except during the surprise task. The robot had a simple state machine for the task-level planning. The human pilot(s) would provide a few clicks in order to seed otherwise autonomous perception algorithms, and then would watch the state machine execute. The user interface showed the trajectories that the robot was about to execute before it actually executed them, allowing the operator to prevent any motion that could have had unintended consequences. Otherwise the robot was operating almost entirely on its own.

## 4.0 RESULTS AND DISCUSSION

### 4.1 DRC FINALS COMPETITION

The research results summarized above were all choreographed in a complex dance that balanced research vs competition-readiness. The unfortunate truth was that the competition rewarded very simple robot strategies that were extensively tested/debugged experimentally. The performer had to balance pushing forward new capabilities with fielding limited capabilities robustly. Most teams chose the more conservative approach, the rules changes to accommodate the weaker technology right up until the end.

MIT's performance in the competition was adequate. They put themselves in a position to win, with an advanced set of capabilities based on autonomy that could have allowed our robot to address the harder versions of the competition rules. This was showcased somewhat during the DRC Finals, after our robot fell and broke its right arm. Most testing was performed using only the right arm, and the performers had planned to use it almost exclusively for all of the manipulation tasks. After the damage during the start of the run, the performers were able to flip a bit in the software, and watch as our robot very seemlessly adapted to using the left arm for all the remaining tasks. The drill task unfortunately required two hand to operate. This change rippled through the system in many subtle ways. An interesting observation of this was the robot's decision to approach the valve task from the right side. This was due to the fact that the robot knew it would need to have room to allow the left hand space to manipulate the valve. This dynamic reasoning was significant as during the practice runs it had always approached the task from the opposite side when the right hand was functioning properly.

Detailed summaries of our performance during the trials are available in [1], and during the finals are available in [17].



*Figure 6 - MIT's Atlas Robot performing driving task at DRC Finals*

The driving task consisted of the robot driving a heavily modified Polaris Ranger through a set of Jersey barriers past a finish line. This required the robot grasping and manipulating a steering wheel while controlling a gas pedal using its foot. The gearing of the Ranger allowed the vehicle to come to a stop without the need to apply pressure to a brake pedal. Additionally, the robot needed an additional mount to allow it to easily dismount the vehicle upon completion. Figure 6 shows the robot completing the driving task at the DRC finals.



Figure 7 - MIT's Atlas robot failing to successfully dismount the Polaris Ranger.

During one of MIT's runs, dismounting the vehicle was not successful. The fall resulting in moderate damage to the robot particularly incapacitating the robot's right arm. The robot was however able to accommodate for the damaged limb and continue accomplishing tasks using the opposite arm.



Figure 8 - Atlas Robot successfully opening a door and crossing the threshold.

The task of opening a door relied heavily on abstraction of plans and conveyance of abstracted surfaces and point clouds. The task also relied on the robot applying only enough force to open the door and not too much force as to send the door flying back into the robots face. The MIT team was able to complete this task and cross the threshold as shown in figure 8.




Figure 9 - A remote view of the valve task as seen by the human operator at base station

One of the more difficult manipulation tasks required the robot to turn a valve. The valve task required substantial surface abstraction given the valve shape and the granularity of the point cloud. This type of surface was similar to the steering wheel in the driving task. Unlike the driving task, however, the robot had to manually grasp the shape without prior operator placement. MIT was successfully able to accomplish the task of grasping and turning the valve.



*Figure 10 - Traversal of the uneven terrain at the DRC finals*

Navigating uneven terrain was an extremely perilous task for the Atlas robots. Unlike some other entries, the Atlas is not equipped to lower its center of gravity enough to minimize balancing issues. Instead, careful planning and execution must be performed. MIT was one of the few teams choosing to perform this task over the alternative of removing debris. The MIT team's research developed over the course of the program allowed the robot to successfully traverse the terrain with no falls.

Finally, the robot was required to complete another balance task of climbing a set of stairs. Negotiating this obstacle required substantial balance and planning to allow the robot to maintain its balance through careful monitoring and adjustment of its center of gravity. The MIT team was able to complete this task and was awarded full points for its completion.



Figure 11 - The robot successfully traversing the stairs obstacle

## 5.0 CONCLUSIONS

Massachusetts Institute of Technology's team developed novel technologies capable of significantly advancing the state-of-the-art in bipedal navigation in dynamic communications degraded environments. Their teams mediocre finish should not undermine the significant advanced made in the fields of bipedal locomotion, stereoscopic visual estimation, state-estimation and planning. The substantial impacts of their research will continue to revolutionize the field of robotics for years to come. The performers published substantial volumes in the form of peer-reviewed conference and journal publications and additionally dissertations and theses on the research areas were published.

## 6.0 REFERENCES

PDFs and supplementary videos for all of these references are available at  
<http://groups.csail.mit.edu/locomotion/pubs.shtml>.

Relatively mature implementations of all of the above algorithms and associated simulations have been made available in our open-source software project Drake at: <http://drake.mit.edu>.

1. R. Tedrake, M. Fallon, S. Karumanchi, S. Kuindersma, M. Antone, T. Schneider, T. Howard, M. Walter, H. Dai, R. Deits, M. Fleder, D. Fourie, R. Hammoud, S. Hemachandra, P. Ilardi, C. Perez-D'Arpino, S. Pillai, A. Valenzuela, C. Cantu, C. Dolan, I. Evans, S. Jorgensen, J. Kristeller, J. A. Shah, K. Iagnemma, and S. Teller. A summary of team MIT's approach to the virtual robotics challenge. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, page 1, Hong Kong, China, May 2014.
2. Scott Kuindersma, Frank Permenter, and Russ Tedrake. An efficiently solvable quadratic program for stabilizing dynamic locomotion. In *Proceedings of the International Conference on Robotics and Automation*, Hong Kong, China, May 2014. IEEE.
3. Maurice Fallon, Scott Kuindersma, Sisir Karumanchi, Matthew Antone, Toby Schneider, Hongkai Dai, Claudia Pérez D'Arpino, Robin Deits, Matt DiCicco, Dehann Fourie, Twan Koolen, Pat Marion, Michael Posa, Andrés Valenzuela, Kuan-Ting Yu, Julie Shah, Karl Iagnemma, Russ Tedrake, and Seth Teller. An architecture for online affordance-based perception and whole-body planning. *Journal of Field Robotics*, 32(2):229-254, September 2014.
4. Robin L H Deits and Russ Tedrake. Computing large convex regions of obstacle-free space through semidefinite programming. In *Proceedings of the Eleventh International Workshop on the Algorithmic Foundations of Robotics (WAFR 2014)*, Istanbul, 2014.
5. Anirudha Majumdar, Amir Ali Ahmadi, and Russ Tedrake. Control and verification of high-dimensional systems with DSOS and SDSOS programming. In *Proceedings of the 53rd Conference on Decision and Control (CDC)*, 2014.
6. Hongkai Dai, Andrés Valenzuela, and Russ Tedrake. Whole-body motion planning with centroidal dynamics and full kinematics. *IEEE-RAS International Conference on Humanoid Robots*, 2014.
7. Maurice F. Fallon, Matthew Antone, Nicholas Roy, and Seth Teller. Drift-free humanoid state estimation fusing kinematic, inertial and LIDAR sensing. In *International Conference on Humanoid Robots*, Madrid, Spain, 2014.
8. Robin Deits and Russ Tedrake. Footstep planning on uneven terrain with mixed-integer convex optimization. In *Proceedings of the 2014 IEEE/RAS International Conference on Humanoid Robots (Humanoids 2014)*, Madrid, Spain, 2014.
9. Robin Deits and Russ Tedrake. Efficient mixed-integer planning for UAVs in cluttered environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, 2015.
10. Hongkai Dai, Anirudha Majumdar, and Russ Tedrake. Synthesis and optimization of force closure grasps via sequential semidefinite programming. In *International Symposium on Robotics Research*, 2015.

11. Russ Tedrake, Scott Kuindersma, Robin Deits, and Kanako Miura. A closed-form solution for real-time ZMP gait generation and feedback stabilization. In *Proceedings of the International Conference on Humanoid Robotics*, Seoul, Korea, November 2015.
12. Maurice F. Fallon, Pat Marion, Robin Deits, Thomas Whelan, Matthew Antone, John McDonald, and Russ Tedrake. Continuous humanoid locomotion over uneven terrain using stereo fusion. In *Proceedings of the 2015 IEEE-RAS International Conference on Humanoid Robotics*, 2015.
13. Scott Kuindersma, Robin Deits, Maurice Fallon, Andrés Valenzuela, Hongkai Dai, Frank Permenter, Twan Koolen, Pat Marion, and Russ Tedrake. Optimization-based locomotion planning, estimation, and control design for the Atlas humanoid robot. *Autonomous Robots*, 40(3):429-455, 2016.
14. Benoit Landry, Robin Deits, Peter R. Florence, and Russ Tedrake. Aggressive quadrotor flight through cluttered environments using mixed integer programming. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, May 2016.
15. Michael Posa, Scott Kuindersma, and Russ Tedrake. Optimization and stabilization of trajectories for constrained dynamical systems. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 1366-1373, Stockholm, Sweden, May 2016.
16. Pierre-Brice Wieber, Russ Tedrake, and Scott Kuindersma. Modeling and control of legged robots. In Bruno Siciliano and Oussama Khatib, editors, *Springer Handbook of Robotics, 2nd Ed*, chapter 48. Springer, 2015.
17. Pat Marion, Maurice Fallon, Robin Deits, Andrés Valenzuela, Claudia Pérez D'Arpino, Greg Izatt, Lucas Manuelli, Matt Antone, Hongkai Dai, Twan Koolen, John Carter, Scott Kuindersma, and Russ Tedrake. Director: A user interface designed for robot operation with shared autonomy. *Journal of Field Robotics*, 1556-4967, 2016.
18. Lucas Manuelli and Russ Tedrake. Localizing external contact using proprioceptive sensors: The contact particle filter. In *International Conference on Intelligent Robots and Systems (IROS)*, October 2016.
19. Twan Koolen, Michael Posa, and Russ Tedrake. Balance control using center of mass height variation: limitations imposed by unilateral contact. In *Humanoid Robots (Humanoids), 2016 IEEE-RAS 16th International Conference on*, pages 8-15. IEEE, 2016.
20. Hongkai Dai and Russ Tedrake. Planning robust walking motion on uneven terrain via convex optimization. *IEEE-RAS International Conference on Humanoid Robots*, 2016.
21. Andrés K. Valenzuela. *Mixed-Integer Convex Optimization for Planning Aggressive Motions of Legged Robots Over Rough Terrain*. PhD thesis, Massachusetts Institute of Technology, Feb 2016.
22. Hongkai Dai. *Robust multi-contact dynamical motion planning using contact wrench set*. PhD thesis, Massachusetts Institute of Technology, 2016.

## **LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS**

DARPA	Defense Advanced Research Projects Agency
DRC	DARPA Robotics Challenge
ICP	Instantaneous Capture Point
IMU	Inertial Measurement Unit
LIDAR	Light Detection and Ranging
MIT	Massachusetts Institute of Technology
QP	Quadratic Program
VTK	Video Toolkit
ZMP	Zero Moment Point